

Searching within **The ACM Digital Library** for: oltp and single pass and grouping change data (str)
Found 29 of 257,584

REFINE YOUR SEARCH[Search Results](#)[Related Journals](#)[Related Magazines](#)[Related Sites](#)**▼ Refine by Keywords**

oltp and single pass and grouping change data

Discovered Terms

▼ Refine by People

Names

Institutions

Authors

▼ Refine by Publications

Publication Year

Publication Names

ACM Publications

All Publications

Content Formats

Publishers

▼ Refine by Conferences

Sponsors

Events

Proceeding Series

ADVANCED SEARCH[Advanced Search](#)**FEEDBACK**
 Please provide us with feedback

Found 29 of 257,584

[Search Results](#)[Related Journals](#)[Related Magazines](#)[Related Sites](#)Sort by

Results 1 - 20 of 29

 Save results to a Binder**1 Cache-oblivious databases: Limitations and opportunities** Bingsheng He, Qiong LuoJune 2008 **Transactions on Database Systems (TODS)**, Volume 33 Issue**Publisher:** ACM Full text available:  Pdf (1.09 MB) Additional Information: [full citation](#), [abstract](#), [bibliometrics](#)**Bibliometrics:** Downloads (6 Weeks): 35, Downloads (12 Months): 316, Citation

Cache-oblivious techniques, proposed in the theory community, have optimal performance in terms of cache misses. They achieve this by avoiding amount of data transferred between any two adjacent levels of an arbitrated hierarchy. This optimal performance is achieved without any ...

Keywords: Cache-oblivious, cache-conscious, chip multiprocessors, data locality, database systems, memory hierarchy, multithreading

2 Computation spreading: employing hardware migration to specialize Koushik Chakraborty, Philip M. Wells, Gurindar S. SohiNovember 2006 **ASPLOS-XII: Proceedings of the 12th international conference on Architectural support for programming languages and operating systems****Publisher:** ACM Full text available:  Pdf (232.44 KB) Additional Information: [full citation](#), [abstract](#), [bibliometrics](#)**Bibliometrics:** Downloads (6 Weeks): 14, Downloads (12 Months): 118, Citation

In canonical parallel processing, the operating system (OS) assigns a processor to a thread from a multithreaded server application. Since different threads from the same application may require similar computation, albeit at different ...

Keywords: cache locality, dynamic specialization

Also published in:

October 2006 **SIGARCH Computer Architecture News** Volume 34 Issue 5October 2006 **SIGOPS Operating Systems Review** Volume 40 Issue 5November 2006 **SIGPLAN Notices** Volume 41 Issue 11**3 Pipeline spectroscopy** Thomas R. Puzak, A. Hartstein, P. G. Emma, V. Srinivasan, Arthur Nadas, June 2007 **ExpCS '07: Proceedings of the 2007 workshop on Experimental computer science****Publisher:** ACMFull text available:  Pdf (305.12 KB) Additional Information: [full citation](#), [abstract](#), [bibliometrics](#)**Bibliometrics:** Downloads (6 Weeks): 6, Downloads (12 Months): 19, Citation

Pipeline Spectroscopy is a new technique that allows us to measure the

of a miss is displayed (graphed) as a histogram, which represents a pre visualization of the cost of each cache miss ...

Keywords: cache, convex combination, cost of a miss, probability tran:

4 DHIS: discriminating hierarchical storage

Chaitanya Yalamanchili, Kiron Vijayasankar, Erez Zadok, Gopalan Sivathan May 2009 **SYSTOR '09: Proceedings of SYSTOR 2009: The Israeli Experi**

Publisher: ACM [Request Permissions](#)

Full text available: [Pdf](#) (208.96 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 8, Downloads (12 Months): 36, Citation

A typical storage hierarchy comprises of components with varying performance providing multiple options for data placement. We propose and evaluate DHIS, that uses application-level hints to discriminate ...

Keywords: file systems, intelligent disks, storage stack, storage system

5 Peta-scale data warehousing at Yahoo!

Mona Ahuja, Cheng Che Chen, Ravi Gottapu, Jörg Hallmann, Waqar Hasan, Kozyrczak, Ramesh Pabbati, Neeta Pandit, Sreenivasulu Pokuri, Krishna Up June 2009 **SIGMOD '09: Proceedings of the 35th SIGMOD international conference on Management of data**

Publisher: ACM [Request Permissions](#)

Full text available: [Pdf](#) (823.92 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 100, Downloads (12 Months): 187, Citation

Insights based on detailed data on consumer behavior, product performance are driving innovation and competition in the internet space. We introduce a warehousing engine, based on a column architecture ...

Keywords: analytics, business intelligence, column database, column store, database, vector query processing

6 Extending query rewriting techniques for fine-grained access control

Shariq Rizvi, Alberto Mendelzon, S. Sudarshan, Prasan Roy June 2004 **SIGMOD '04: Proceedings of the 2004 ACM SIGMOD international conference on Management of data**

Publisher: ACM [Request Permissions](#)

Full text available: [Pdf](#) (172.57 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 16, Downloads (12 Months): 122, Citation

Current day database applications, with large numbers of users, require mechanisms, at the level of individual tuples, not just entire relations/views, so that data can be accessed by each user. Fine-grained ...

7 Replay-based approaches to revision processing in stream query engines

Anurag S. Maskey, Mitch Cherniack March 2008 **SSPS '08: Proceedings of the 2nd international workshop on Stream processing and stream systems**

Publisher: ACM

Full text available: [Pdf](#) (433.78 KB)

Additional Information: [full citation](#), [abstract](#),

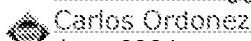
Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 53, Citation

Data stream processing systems have become ubiquitous in academic a

application areas that include financial services, network traffic analysis control. The append-only model of streams implies ...

Keywords: corrections, revision processing, revisions, streams

8 Horizontal aggregations for building tabular data sets



June 2004 **DMKD '04:** Proceedings of the 9th ACM SIGMOD workshop on and knowledge discovery

Publisher: ACM

Full text available: [Pdf](#) (183.18 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 53, Citation

In a data mining project, a significant portion of time is devoted to build analysis. In a relational database environment, building such data set u aggregating columns with SQL queries. Existing ...

9 Bridging the gap between OLAP and SQL

Jens-Peter Dittrich, Donald Kossmann, Alexander Kreutz

August 2005 **VLDB '05:** Proceedings of the 31st international conference or

Publisher: VLDB Endowment

Full text available: [Pdf](#) (409.18 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 90, Citation

In the last ten years, database vendors have invested heavily in order t features for decision support. Examples of functionality that has been a [7], spreadsheet computations [19], grouping ...

10 SQL memory management in Oracle9i

Benoit Dageville, Mohamed Zait

August 2002 **VLDB '02:** Proceedings of the 28th international conference or

Publisher: VLDB Endowment

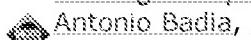
Full text available: [Pdf](#) (107.63 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 4, Downloads (12 Months): 52, Citation

Complex database queries require the use of memory-intensive operato operators need memory, also referred to as SQL memory, to process th sort operator uses a work area to perform the in-memory ...

11 Adding subqueries to MySQL, what does it take to have a decision-s



November 2002 **DOLAP '02:** Proceedings of the 5th ACM international work OLAP

Publisher: ACM

Full text available: [Pdf](#) (200.24 KB)

Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 11, Downloads (12 Months): 83, Citation

MySQL is an OLTP system without the ability to handle SQL subqueries. optimization, we add subquery processing to MySQL through rewriting. handle subqueries in the FROM clause; then ...

Keywords: SLQ, query optimization, unnesting

12 Queue: Volume 6 Issue 3

 May 2008 Queue
Publisher: ACM

Full text available:  Digital Edition ,  Pdf (5.56 MB) Additional Information: [full citation](#)

Bibliometrics: Downloads (6 Weeks): 107, Downloads (12 Months): 723, Citations: 11

13 Compression techniques for fast external sorting

 John Yiannis, Justin Zobel

April 2007 **The VLDB Journal — The International Journal on Very Large Databases**, Volume 21 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available:  Pdf (433.12 KB) Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): 13, Downloads (12 Months): 117, Citations: 11

External sorting of large files of records involves use of disk space to store the sorted data, time for sorting, and transfer time between CPU, cache, memory, and disk. We propose compression techniques to reduce the storage and transfer costs, and, in the case of external ...

Keywords: External sorting, Query evaluation, Semi-static compression, Data compression

14 Communications of the ACM: Volume 52 Issue 9

 September 2009 **Communications of the ACM**

Publisher: ACM

Full text available:  Digital Edition ,  Pdf (7.28 MB) Additional Information: [full citation](#)

Bibliometrics: Downloads (6 Weeks): 2250, Downloads (12 Months): 2250, Citations: 22

15 Compiling mappings to bridge applications and databases

 Sergey Melnik, Atul Adya, Philip A. Bernstein

November 2008 **Transactions on Database Systems (TODS)**, Volume 33 Issue 4

Publisher: ACM 

Full text available:  Pdf (1.40 MB) Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): 38, Downloads (12 Months): 455, Citations: 45

Translating data and data access operations between applications and databases is a management problem. We present a novel approach to this problem, in which the application data and the persistent storage is ...

Keywords: Mapping, query rewriting, updateable views

16 PNUTS: Yahoo!'s hosted data serving platform

Brian F. Cooper, Raghu Ramakrishnan, Utkarsh Srivastava, Adam Silberstein, Jacobson, Nick Puz, Daniel Weaver, Ramana Yeremni

August 2008 **Proceedings of the VLDB Endowment**, Volume 1 Issue 2

Publisher: VLDB Endowment

Full text available:  Pdf (495.87 KB) Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): 26, Downloads (12 Months): 224, Citations: 22

We describe PNUTS, a massively parallel and geographically distributed system for serving web applications. PNUTS provides data storage organized as hashed or replicated buckets, and handles large numbers of concurrent requests including updates ...

17 Retrospective on Aurora

Hari Balakrishnan, Magdalena Balazinska, Don Carney, Uğur Çetintemel, Michael Chernikov, Eddie Galvez, Jon Salz, Michael Stonebraker, Nesime Tatbul, Richard Tibbe

December 2004 **The VLDB Journal — The International Journal on Very Large Data Bases**
Issue 4

Publisher: Springer-Verlag New York, Inc.

Full text available:  [Pdf \(349.43 KB\)](#) Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 47, Citation

This experience paper summarizes the key lessons we learned through the development of the Aurora stream-processing engine. For the past 2 years, we have developed many applications using Aurora. We first describe in detail ...

Keywords: Data stream management, Distributed stream processing, Internet-of-service, Stream-processing engines

18 MIL primitives for querying a fragmented world

Peter A. Boncz, Martin L. Kersten

October 1999 **The VLDB Journal — The International Journal on Very Large Data Bases**
Volume 13 Number 2

Publisher: Springer-Verlag New York, Inc.

Full text available:  [Pdf \(261.36 KB\)](#) Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 37, Citation

In query-intensive database application areas, like decision support and data warehousing, vertical fragmentation have a significant performance advantage. In order to build efficient query engines for such applications on top of such a fragmented ...

Keywords: Database systems, Main-memory techniques, Query languages, Vertical fragmentation

19 Database-aware semantically-smart storage

Muthian Sivathanu, Lakshmi N. Bairavasundaram, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau

December 2005 **FAST'05: Proceedings of the 4th conference on USENIX Annual Technical Conference**
Storage Technologies - Volume 4, Volume 4

Publisher: USENIX Association

Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citation

Recent research has demonstrated the potential benefits of building storage systems on top of file systems above them. Such "semantically-smart" disk systems use knowledge structures and operations to improve performance, availability, ...

20 To tune or not to tune?: a lightweight physical design adviser

Nicolas Bruno, Surajit Chaudhuri

September 2006 **VLDB '06: Proceedings of the 32nd international conference on Very Large Data Bases**

Publisher: VLDB Endowment

Full text available:  [Pdf \(733.90 KB\)](#) Additional Information: [full citation](#), [abstract](#),

Bibliometrics: Downloads (6 Weeks): 7, Downloads (12 Months): 44, Citation

In recent years there has been considerable research on automating the physical design of DBMSs. Current techniques provide good recommendations, but are relatively conservative when deciding to launch a resource-inten-

